Automated Planning for the Modified Antarctic Mapping Mission[1]

Benjamin D. Smith, Barbara E. Engelhardt, Darren H. Mutz, John P. Crawford

Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
M/S 126-347
Pasadena, CA 91109
~~(818) 393-5371~~
{firstname.lastname}@jpl.nasa.gov

Abstract—~~—~~The RadarSAT Modified Antarctic Mapping Mission (MAMM) ran from September to November 2000. The MAMM mission consisted of over 2400 synthetic aperture radar (SAR) data takes over Antarctica that had to satisfy coverage and other scientific criteria while obeying tight resource and operational constraints. Developing these plans is a time and knowledge intensive effort. It required over a work-year to manually develop a comparable plan for AMM-1, the precursor mission to MAMM. This paper describes the automated mission planning system ~~The ASPEN planning system automated~~ for MAMM, which dramatically reduced mission planning costs to just a few work-weeks, and~~the mission planning process~~ enabled rapid generation of "what-if" scenarios for evaluating mission-design trades. This latter capability informed several critical design decisions and was instrumental in accurately costing the mission. ~~and provided a fast replanning capability for responding to anomalies during operations. Furthermore, the ASPEN planning system enabled mission analysis and what-if scenarios early on in mission design, which significantly impact critical mission design decisions.~~ This paper describes the mission, the planning problem, the system architecture, the planning challenges involved, and ~~how the~~the impact of the automated planning system on planning and operating the mission.~~ ASPEN planning system helped in mission design and operations.~~

TABLE OF CONTENTS

1. INTRODUCTION

The Modified Antarctic Mapping Mission (MAMM) executed from September through November of 2000 onboard RadarSAT, a Canadian Space Agency (CSA) satellite. This joint NASA/CSA mission is a modified version of the First RadarSAT Antarctic Mapping Mission (AMM-1) executed in 1997. The objective of AMM-1 was to acquire complete coverage of the Antarctic continent, whereas the objective of MAMM is to acquire repeat-pass interferometry to measure ice surface velocity of the outer regions of the continent, north of 80S. The mission objective is to perform synthetic aperture radar (SAR) mapping of the Antarctic over three consecutive 24-day repeat cycles. The SAR instrument has several "beams" each of which can be commanded to take data in rectangular swaths. These swaths are eventually compiled into a mosaic. The incidence angle of each beam is separated by a few degrees and partially overlaps the swaths of adjacent beams. The location of the swaths at any given time is determined by the spacecraft orbit. The planning problem is to select a subset of the available swaths that fully cover the visible area of Antarctica and satisfy operational and resource constraints imposed by the RadarSAT Mission Management Office (MMO). The driving operational constraints are the limited on-board tape recorder (OBR) capacity and downlink opportunities, which constrain the swath subsets that will fit on the OBR between downlinks.

The AMM-1 ~~experience~~ mission demonstrated the need for an automated planning capability. ~~at manually developing mission plans was laborious and error-prone.~~ The schedule for AMM-1 consisted of 850 acquisitions (swaths) over 18 days, and took over a work-year to develop manually. Despite repeated checking, t~~The plans took months to develop, and~~ his plan violated operations ~~some~~ constraints that ~~violations~~ were not detected until the final MMO

[1] 0-7803-6599-2/01/$10.00 © 2001 IEEE

review. This inability to detect all the operations and resource constraint violations during the planning process required expensive and disruptive last-minute revisions. An automated planner could have quickly identified constraint violations, suggested repairs, and reduced the chance of errors, all of which would have significantly expedited the mission planning process.

This experience led to the development and use of an automated mission planning system for MAMM. The system expanded a set of swaths selected by the human mission planner into a detailed plan, automatically scheduled downlink activities to minimize resource costs and other criteria, and checked the resulting plan for operations constraint violations. With this system MAMM developed a 24-day mission plan containing 800 swaths in a matter of weeks, as compared to the work-year required to develop a comparable mission plan for AMM-1.

In addition to reducing the plan development effort, the MAMM planner also provided resource tracking and other details that enabled accurate costing and feasibility estimates. The MAMM planner also enabled "what-if" studies that were not possible under AMM-1. The planner quickly generated detailed variations of the baseline plan for different ground station availability assumptions. These study plans were instrumental in selecting ground stations and making other decisions about mission alternatives.

~~Anomalies during AMM-1 operations caused many data takes to be lost. The missing data had to be rescheduled for later in the mission. Any changes to the plan had to be submitted within 36 hours of any reacquisition, which meant replan options had to be identified as quickly as possible in case the optimal changes to the schedule were within 36 hours of the anomaly identification. To manually turn around plans within these time constraints required a team of four people working from pre-generated contingency plan segments. The missed observations were placed into gaps in the original plan to minimize coverage holes. More extensive changes, such as altering the original planned swaths or attempting to put all of the real-time observations on the tape were avoided to minimize the planning effort and the chance of introducing errors into the plan. Holes in coverage that could not be replanned trivially were simply ignored, creating significant gaps in coverage. Automated replanning during operations would allow faster turn-around with fewer people, and enable more extensive changes to the schedule in order to maximize science return.~~

The rest of this paper describes the automated planning system that was constructed for MAMM based on the ASPEN [1,4,5] planning environment. ~~This system developed baseline and contingency mission plans and was a part of the replanning scenario during operations to reschedule images missed due to anomalies.~~Section 2 describes the mission planning problem, Section 3 describes the automated planning system, Section 4 describes the impact the system had on mission planning and operations, and future work and conclusions appear in Section 5.

2. ~~2.~~ MAMM Planning Problem ~~Description~~

The objective of MAMM is to acquire repeat-pass SAR interferometry of Antarctica north of –80 degrees latitude over three consecutive 24-day repeat cycles[2] to measure ice surface velocity of the outer regions of the continent.

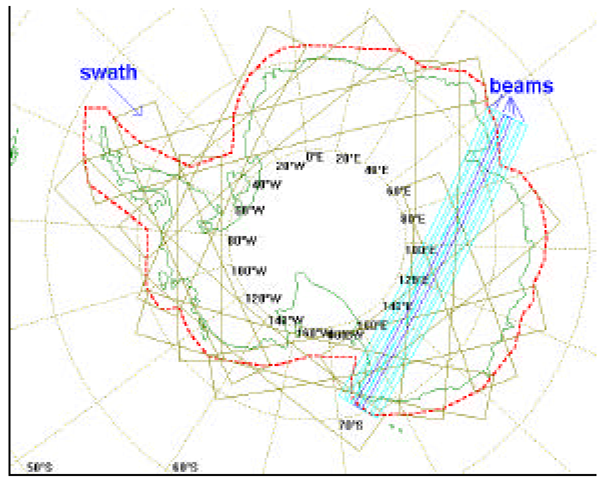

Figure 1: Swath Selection Problem. The mission planne must select swaths and beams that meet coverage and scientific criteria, and also satisfy operations constraint

MAMM will use fine beams (high resolution) from RadarSAT to increase the accuracy of the interferometric data analysis for fast-moving glaciers found in the AMM-1 mission.

Interferometry requires multiple images, each taken with the same beam and on the same relative orbit within a repeat cycle.~~The MAMM mission is to acquire full coverage of all of visible Antarctica three times in order to perform repeat-pass inferometry to map the surface velocity field. Modified from AMM-~~

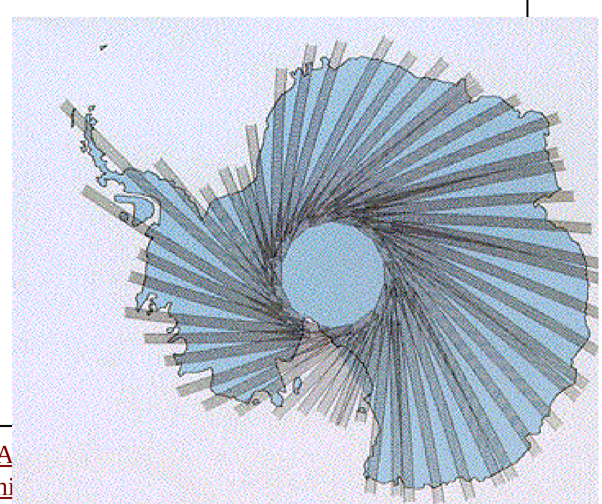

Figure 2: Map of Antarctica showing a single bea coverage map (this is not entirely accurate as fi beams will be used for most of the costal coverag Notice the coverage gap in the center due to lack

[2] The RadarSA
306 orbits, whi
cycle inscribes
other cycle.

~~1, the satellite will remain in north looking mode during MAMM, restricting coverage to regions north of about –80 degrees Latitude. MAMM will use fine beams from RadarSAT to increase the accuracy of the inferometric data analysis for fast-moving glaciers found in the AMM-1 mission.~~ The mission plan consists of SAR acquisitions and data downlinks for one 24-day repeat cycle. This plan is repeated over a total of three consecutive repeat cycles in order to obtain identical beam/orbit images taken on consecutive cycles as required for interferometry. Since the acquisitions for all three cycles must be the same, data takes lost to anomalies in the second and third cycles will not be rescheduled. However, data takes lost ~~Inferometry requires multiple images, each with the spacecraft on identical orbit and using identical beams, and moreover requires both ascending and descending passes (where *descending* is imagery taken when the spacecraft trajectory is headed toward the south pole and *ascending* is when the spacecraft trajectory is headed away from the south pole).~~ during the first cycle will be rescheduled for later in that cycle, and the modified schedule will be executed on cycles two and three.

The Planning Process

The mission planning process consists of the following four steps:

1. Select SAR swaths that cover the desired target regions in Antarctica and satisfy other scientific requirements. Each target must be imaged once on an ascending pass, and once on a descending pass. The swaths are selected from all the swaths that intersect the target regions during one 24 day repeat cycle. This problem is illustrated in Figure 1.
2. Expand the selected swaths into a complete schedule. Acquiring a SAR image (swath) can involve turning on the instrument, changing beams, spinning up the recorder, and other actions. Then determine whether the expanded schedule violates any operations constraints (e.g., there is insufficient time to make a beam-switch between two data-takes that require different beams).
3. Create a downlink schedule. Each image must either be downlinked in real-time to a ground station that is in view during the acquisition, or stored to the data recorder and downlinked at a later opportunity. The downlink schedule must obey resource and operations constraints (e.g., recorder capacity, station visibility, time for station lock-up and exit) and conform to a priority policy (certain stations are more reliable or have lower costs than others; resource costs make real-time takes preferable to recorded ones).
4. If the schedule violates operations constraints, or the downlink schedule is of insufficient quality or cannot transmit all of the acquisitions, return to Step 1 and modify the selected swaths to correct the problems. Modifications include changing the swath start time, duration, and/or beam; or selecting an alternate swath on a different orbit that covers the same target area.

This process results in an acquisition plan. Acquisitions lost during execution of the first cycle are rescheduled using a similar process—identify an alternate swath that covers the missed target region, perform Steps 2 and 3 to see if the resulting schedule is valid, and if not return to Step 1 and modify the swath or select another candidate. The other swaths in the plan are not modified in order to minimize disruption.

The MAMM planning system automates Step 2 and 3, and reports the information needed for Step 4: operations violations, data that cannot be downlinked, and schedule quality summary. It is up to the mission planner to select (Step 1) and modify (Step 4) the SAR swaths, since these steps require human judgement of the science impact. Swath selection (Step 1) is partially automated by a tool called SPA [7], developed by the Canadian Space Agency, that identifies the available swaths by propagating the spacecraft orbit but does not consider operations or downlink constraints. The future work section discusseses how automated planning technology could further automate Steps 1 and 4 with varying degrees of input from human experts.

The Planning Problem

The MAMM and AMM-1 missions were conducted aboard RadarSAT. For purposes of planning, the satellite can perform the following activities:

- Perform a SAR data take and save it to the onboard recorder (OBR Take)
- Perform a SAR data take and downlink it in real time (RTM Take)
- Switch the SAR beam
- Start the tape recorder (takes a few seconds to spin-up)
- Stop the tape recorder (takes a few seconds to spin-down)
- Playback the entire tape and downlink it.

- Establish a downlink session with a ground station.

A mission plan consists of a list of activities, each with a fixed start time and duration. These activities consume spacecraft resources, which are also tracked in the plan. The relevant resources are the tape capacity used, instrument on time per orbit, and number of tape start/stop cycles. In planning parlance, these are all *depletable* resources. The first two of these depletable resources are *renewable*: the tape capacity is replenished by playback activities, and the instrument on-time is replenished once per orbit.

The mission plan must obey certain operational constraints. For example, recording acquired data to the OBR must occur while the tape recorder is running. OBR playbacks and real time (RTM) acquisitions must occur during an established downlink session. There are a total of 24 operational constraints, some of which are shown in Table 1.

There are four ground stations that can potentially receive RadarSAT data for MAMM. A downlink session can only be established with one station at a time, and only while the station is *in view* of the satellite. These requirements are specified as a list of downlink opportunities, or *masks*, each of which consists of a start time, duration, and station identifier. A downlink session can only be established with a station when the satellite ground track is within one of these station masks. The opportunities vary in duration, but are always shorter than the tape capacity. A ground station can acquire RTM data in parallel with an OBR playback.

The planning problem primarily consists of selecting swaths and downlink opportunities. The key decisions are which swaths to select, whether to take a given swath in OBR or RTM mode, and at which downlink opportunity to play back data. In making these decisions, the planner (human or automated) must not only meet hard constraints, but also try to maximize preference criteria. These include preferring ground stations that have higher reliability or lower cost, preferring RTM data to OBR (to minimize ground station and OBR resource costs), and selecting swaths to maximize coverage and other scientific criteria.

The swath and downlink selection decisions are tightly couple. Minor alternations to one part of the schedule tend to require a cascade of additional changes throughout the schedule, which is part of what makes this planning problem difficult. For example, changing a swath from RTM to OBR will increase the tape usage, which may invalidate a previously selected downlink opportunity (if its duration is shorter than the recorded data volume to be downlinked). A different (longer) downlink opportunity must be chosen, and that choice may require changes to other data acquisitions, and those changes may impact other downlinks, and so on throughout the schedule.

As another example, the temporal spacing between adjacent data takes determines whether the tape recorder keeps running between takes or can be stopped, which in turn impacts the tape usage. Changing the temporal spacing between data takes can therefore change the tape usage,

Table 1: Selected Operations Constraints

OBR data can only be downlinked when a ground station outer (or inner) mask is in view
Gap between last data take and transmit must be > 2*(tape_length - tape_remaining)
All data must be downlinked
OBR tapes cannot playback outside of the outer mask
Two activities using SAR cannot happen simultaneously
Cannot transmit RTM data when recorder is in record, spin-up, or spin-down modes
Data takes shall be no less than 1.0m (including 8s pads)
Data takes shall be at least 5.25s apart when beams are changed
Data takes shall be at least 11s apart when beams are not changed
There will be a maximum of 6 OBR transactions per orbit
OBR takes 10s to spin up, consumes 10s of tape
OBR takes 5.5s to spin down, and consumes 5.5s of tape.
OBR spin-up/spin-down between takes iff OBR data takes are > 30s apart
To downlink OBR data, OBR must be in playback mode
OBR cannot record during playback or record during RTM data take
SAR shall be on at most 32.0 minutes per orbit

which may have cascading effects through the schedule.

~~There are 3 steps to create a complete coverage map, all of which are constrained by spacecraft, orbital, and political constraints. (1) Create two complete coverage maps of Antarctica visible from the non-rotated spacecraft, one ascending and one descending, for one spacecraft cycle (containing 306 complete orbits). (2) Create a downlink schedule for each of the images, using either the on-board recorder (limited to 650 minutes of datatakes per cycle) or real-time downlinking (must be to McMurdo ground station, when it is visible at the time of the imaging). (3) If any image is lost due to anomalies during the first cycle, attempt to reschedule a swath later on in that cycle to cover the lost ground and retain complete coverage. The submissions must be within 29 hours of the altered datatake, which is a seven hour improvement over the first mission but does not offer much additional time for planners. But since there are relative orbit repeats every 3.5 days and then every 7 days, a typical replan will not alter swaths for 3 days from the anomaly.~~

~~The next two cycles will use the same schedule as the first cycle, including any replans, so as to obtain complete inferometry data of Antarctica. Anomalies in the second and third cycles will be ignored, because replanning will impact inferometry data.~~

~~Automated planning can help with this process in each of the three steps. These steps are described below.~~

3. ~~3.~~ AUTOMATED PLANNING SYSTEM

The automated planning system takes a list of swaths and downlink opportunities (masks), and produces a detailed plan that assigns swaths to downlink opportunities, tracks resource usage, and reports operations constraint violations.

Specifically, it takes as input a list of SAR swaths selected by the mission planner, a list of accessibility masks for each ground station, the station priority policy, and station capabilities (real-time downlink [RTM] and/or on-board recorder playback [OBR]). The swath input specifies the time, duration, and beam of each swath. The user generates this input using a swath selection and coverage analysis tool called SPA, which CSA developed for RadarSAT missions. The mask files are provided by the RadarSAT Mission Management Office (MMO).~~The first basic planning problem is to select a subset of the available swaths that will cover the Antarctic within the 24-day cycle while satisfying all of the MMO constraints. This requires a combination of constraint reasoning, for which planners are well suited, and geometric reasoning for which special purpose algorithms must be called.~~

~~The planning system can operate in two modes: mixed-initiative and automated. In mixed-initiative mode the human user selects the swaths using a coverage analysis tool (Swath PlAnner (SPA), developed with the Canadian Space Agency [7] or Satellite Tool Kit (STK), an off-the-shelf satellite simulator and orbit propagator developed by Analytical Graphics, Inc. [8]). ASPEN takes this set of swaths along with a list of masks from all of the possible downlink stations, including information on whether the priorities for downlink station usage and whether OBR or RTM data can be transmitted to them. ASPEN expands the swaths into a detailed, time ordered plan. The expansion primarily consists of deciding which of the given downlink opportunities to use (thus partially fulfilling the second step), tracking resource usage to monitor usage violations, and verifying adherence to operational constraints. ASPEN then reports any constraint conflicts that it cannot resolve without modifying the swaths (e.g., they oversubscribe the on-board recorder and cannot be downlinked)~~

The mask and swath files are combined into a single file and passed to the ASPEN planning system, which is described in more detail below. The planner expands the swaths and masks into a detailed plan that includes downlink session activities, tape on/off transitions, beam switches, and tracks resource usage. ASPEN then checks the resulting plan for operations constraints violations. The resulting plan and violations are then converted from ASPEN format to a time-ordered sequence of events and constraint violations in an Excel format that was specified by the mission planners. It also summarizes plan metrics, such as total on-board and ground station resource consumption. This flow of information is documented in Figure 2.

Based on the report files, the human ~~user~~mission planner modifies the ~~swath selection~~selected swaths ~~or downlink station schedule based~~ as needed to resolve the conflicts or improve schedule quality. The check-and-edit cycle is repeated until a conflict-free plan is generated. This rapid feedback allows the user to generate a conflict-free plan much more ~~rapidly~~ quickly than is possible by hand.

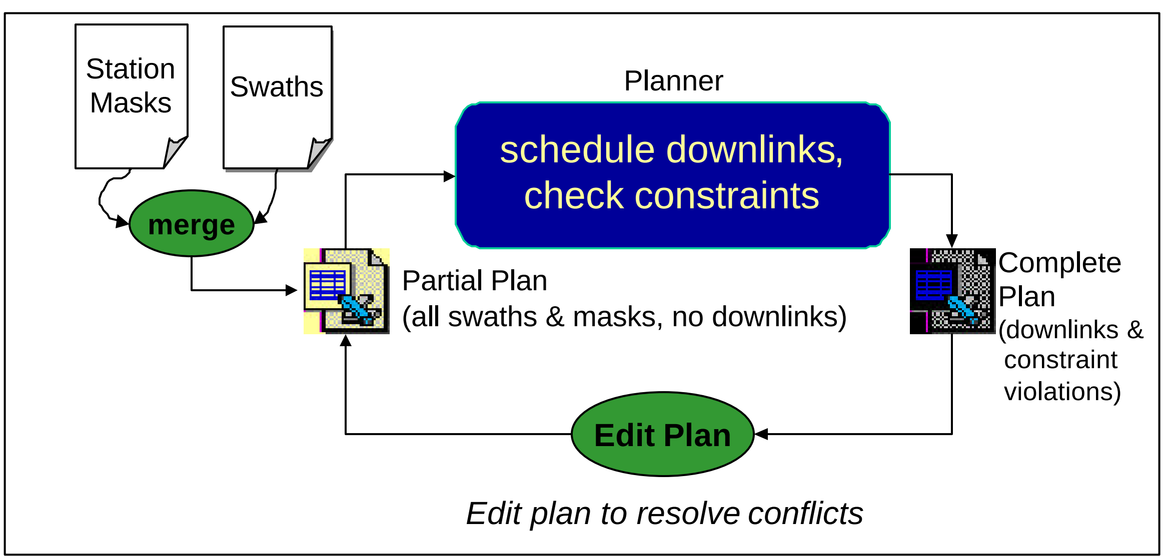

Figure 2: System Data Flow Architecture

~~Moreover, he~~Maintaining the human planner in the loop ~~mixed-initiative mode~~ enables ~~allows the user~~the use of human scientific judgment in selecting~~swaths or downlink selection~~ swaths.~~In automated mode, ASPEN solves the problem automatically. Although this mode was not required as part of the final delivered ASPEN model, these are the specifications for the system. The user provides a set of swath opportunities for each missed observation, based on possibilities generated using the swath analysis tools. ASPEN selects at most one opportunity for each observation such that the resulting plan satisfies the operational constraints and recovers as many missed observations as possible. ASPEN can be forced to use a specific swath to recover a given observation by providing only one opportunity. ASPEN solves the overall planning problem with a combination of forward dispatch and a specialized placement algorithm [4,6]. A specialized set-covering algorithm (e.g., [3]) provides a solution to the swath selection problem, which guides these two algorithms.~~

~~This mode can be used primarily in operations, to accomplish the third objective of replanning, when new plans have to be turned around quickly following anomalies. Anomalies result in complete or partial missing datatakes, which must be rescheduled. The user provides ASPEN with a set of swaths for the missed observations. This same capability could also support a more aggressive mode where ASPEN generates an initial mission plan from scratch, which the user then modifies to meet unarticulated scientific preferences. The user declares that each rectangular region of Antarctica to be imaged is a missed observation, and the swath analysis tools themselves can produce a set of opportunities for each one. ASPEN then selects swaths for these missed observations as it would in anomaly replanning.~~

The ~~4. System Architecture~~

~~The planning system takes as input a set of swath opportunities, downlink opportunities, and a partial plan. The downlink opportunity file is provided by the MMO, and the user generates the swath opportunities file from a coverage analysis tool (SPA). The swath opportunities may be fixed so that there is only one swath for each missed observation (so that ASPEN has no swath selection decisions to make), or open so that ASPEN must select among several opportunities for each missed observation. The partial plan can also force downlink selection decisions, or leave them up to the currently implemented algorithm in ASPEN. This conflict-free plan is then converted into a spreadsheet format, which is the format used by the MAMM mission planners.~~

~~The downlink and partial plan files are converted into an ASPEN plan file. ASPEN expands the partial plan into a complete plan that satisfies the MMO constraints as encoded in the domain model. The model contains external dependency functions that ensure the datatake activities in the plan are consistent with the swaths in the swath opportunities file. The planning algorithm also consults this file to perform swath selection. If there is only one opportunity, the swath selection is trivially solved.~~

~~ASPEN generates a plan file and a list of conflicts that it was unable to resolve (e.g., a swath durations was too short and violated mission or spacecraft constraints). The plan file is converted into an Excel format that the mission planners prefer, and a list of swaths in SPA format (the swath request format required by the MMO). If the swath opportunities file was fixed, this is a pass-through operation; otherwise it is a down-selection of the original file. It will also generate a swath file in Satellite Tool Kit (STK) format. STK has more powerful coverage analysis capabilities than SPA, since it has an orbit propagation tool and can perform coverage analysis. The flow of information is documented in Figure 4, which shows the automated version of this mission planning system.~~ASPEN planner for MAMM

ASPEN [1,4,5] is an automated planning and scheduling system developed at the Jet Propulsion Laboratory and used for a number of space applications. Its basic operation is to find a detailed course of action—or *plan*—that achieves specified high-level goals. The goals, the actions it can take, and the operations constraints on the plan are specified in a declarative *domain model*.

The ASPEN planner has an incremental constraint tracking facility and a search facility. It uses these to process MAMM plans as follows. The search facility generates downlink activities and expands the initial plan into a detailed plan. The constraint tracker determines whether the expanded plan violates any of the constraints in the domain

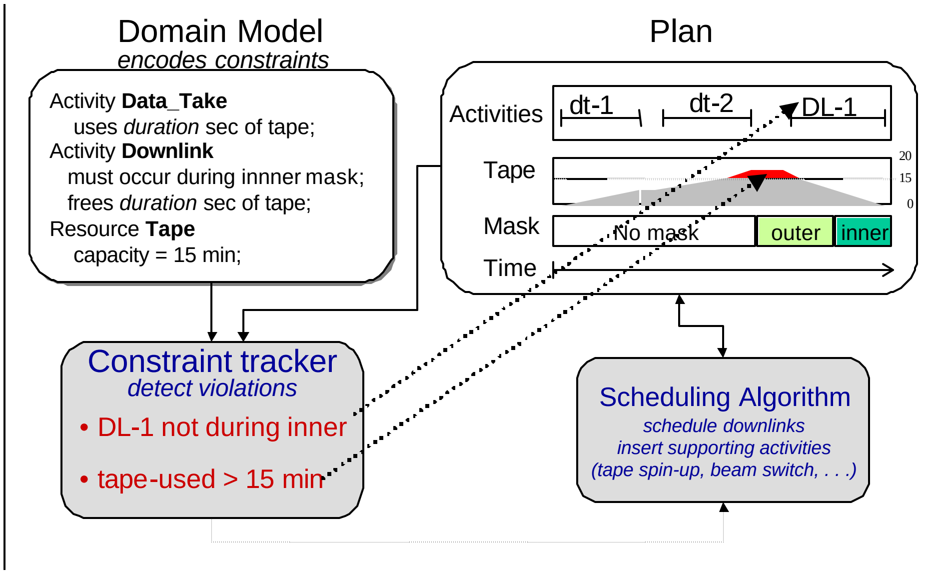

Figure 3: ASPEN planning components

model. This structure is shown in Figure 3. The remainder of the section discusses this operation in more detail.

An ASPEN plan consists of three elements: activities, states, and resources. An activity is an action the spacecraft can perform, such as a data take or beam switch. Activities have a start time and duration and may overlap each other. A resource represents a physical or logical resource of the spacecraft, such as the onboard recorder tape or instrument on-time. A state represents a physical or logical state of the spacecraft, such as the current SAR beam or whether a given ground station is *in-view* or *not- in- view*. Each state and resource is represented as a *timeline* that shows how it evolves over time. Figure 4 shows a sample plan fragment with each of these elements.

The MAMM domain model defines the following major activities:

- DataTake(start, duration, OBR-RTM mode)
- BeamSwitch(from,to)
- SpinUp [start the tape recorder]
- SpinDown [stop the obr]
- OBRContinue [keep running between takes]
- Playback [play the entire tape and downlink it.]
- Downlink [establish a downlink session]
- InView (one for each station)
- OrbitStart

The InView activities populates the station mask timelines, and is not an executable activity. The OrbitStart activity resets the SAR instrument on-time resource to zero (the SAR on-time has a per-orbit maximum).

The MAMM domain model defines these major states and resources:

• SAR_on_time	[resource]
• OBR_storage	[resource]
• OBR_tape_transactions	[resource]
• StationInView (yes,no)	[state] (1 per station)
• Beam (S1-S7,E1-E7)	[state]
• OBR_mode (idle, playback, rec.)	[state]

These plan elements are related by *constraints*. These can be temporal constraints among activities (a tape spin-down must immediately follow a data take), resource constraints (a data take uses *d* seconds of OBR tape, where *d* is the duration of the data take), and state constraints (the SAR instrument must be ON during a data take). The MAMM operations constraints were encoded in terms of these constraints.

The planner's constraint tracking facility maintains the state and resource timelines for the current plan and determines whether the constraints are satisfied. Whenever the plan changes, it incrementally recomputes the impacted timelines and constraints. State and resource timelines are computed from the state and resource constraints imposed by the activities in the current plan.

The planner takes as input a partial plan that contains just the swath activities and the downlink mask activities. The mask activities populate the mask timelines for each ground station. The planner then decides how to downlink the swaths. This downlink assignment phase assigns swaths to downlink opportunities (masks) and records these decisions in the plan by adding downlink activities for each selected mask and grounding the 'downlink mode' parameter of each swath activity to OBR or RTM accordingly.

The downlink scheduling problem is a constrained assignment problem. Each swath must be assigned to exactly one downlink opportunity, and that assignment must satisfy OBR constraints (the duration of the selected opportunity must exceed the amount of recorded OBR data) and operations constraints (can only downlink a data take RTM if it is contained by an RTM capable mask; cannot downlink to two stations simultaneously). This problem is solved by a greedy algorithm. In each iteration it makes the best feasible assignment. If no assignment is possible, it backtracks. Since there may be no way to downlink all the selected swaths, it

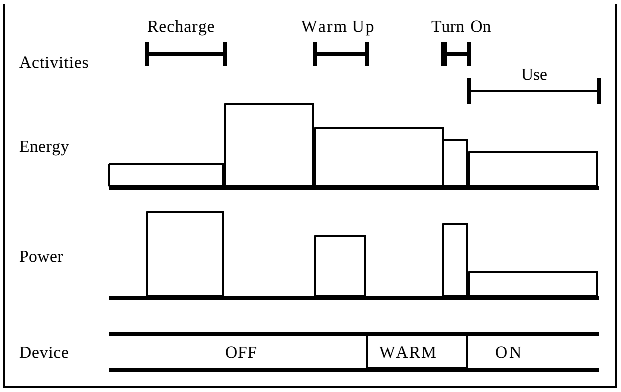

Figure 4: Timelines for activities, two resources (energy and power), and a state variable (device). Each box on the timelines is called a *timeline unit* and represents the value of that state or resource over that time period.

limits its backtracking to a two-orbit window. If no feasible solution can be found in that window, it selects a feasible schedule that downlinks the most data, and reports the lost data as a constraint violation.

After adding downlink activities and grounding the 'downlink mode' parameters of the data take activities the planner performs a limited expansion and grounding of the plan. At this point the plan consists solely of swath, mask, and downlink activities. In each iteration it selects a value for an ungrounded activity parameter, or adds an activity to satisfy an open temporal constraint. For example, if activity A is in the plan and has an open constraint that it must be before activity B, the planner will add an activity instance of type B just after activity A. At the end of this phase, the plan contains all of the activities needed to acquire and downlink the requested swaths. The resource and state timelines have also been computed based on the reservations made by the activities.

For a 24-day plan with 819 swaths and 1,068 downlink masks, the expanded plan contains 8,825 activities and over 16,000 timeline units. In a plan this size the expansion must be performed carefully to avoid unnecessary computation. When an activity is added to the schedule and imposes a resource reservation, it forces all of the resource timeline units downstream of the activity to be recomputed. Placing activities in increasing time order, where possible, minimizes the computation effort.

The expansion uses heuristics to ensure the most computation-efficient ordering. It also uses heuristics in selecting values for grounding parameters, and for resolving disjunctive constraints. For MAMM these choices can always be made correctly based on the surrounding plan context. The expansion heuristics must therefore ensure that those parts of the plan have been expanded before making one of these choices.

One the plan is expanded and grounded, the planner uses its constraint tracking mechanism to identify *conflicts*: violations of constraints in the domain model. These consist of temporal violations (e.g., data take activities are too close together), resource violations (e.g., exceeded tape capacity), and state violations.

Finally, the plan and the conflicts are converted into a spreadsheet format. This is a time-ordered list of swath, mask, and downlink activities, with one row for each activity. There is one column for each resource. The value of that column for each activity (row) is the value of that resource at the end of that activity. The last column holds a list of the operations constraint violations in which that activity is involved. A table maps ASPEN conflicts to corresponding high-level operations constraints, and it is these high-level constraints that are reported in the spreadsheet.

4. ~~5.~~ USING AUTOMATED PLANNING

Before the Mission

The mission planner used the ~~constraint checking capabilities for the mixed-initiative mode are completed and have been used to~~automated planning system to generate the MAMM plan. This consisted of a draft plan, a final plan, and several contingency plans. The MAMM plan consisted of 800 swaths over 24 days, and required approximately three work-weeks to develop. ~~generate mission plans, submitted to the Canadian Space Agency. The planner takes a set of swaths and downlink opportunities as input, assigns the swaths to downlink opportunities, adds supporting activities for each swath and checks the resulting plan for constraint violations.~~ By comparison, the mission plan for ~~These~~the first Antarctic Mapping Mission (~~plans took less than three weeks to develop, whereas the plans manually developed and checked for~~ AMM-1) consisted of 850 swaths over 18 days, and required over a work-year to develop manually.

The first draft plan was generated in under two weeks. This was an iterative process, the bulk of which was spent selecting swaths and resolving conflicts. This draft plan provided valuable information for feasibility analyses, and for first order estimates of ground station costs and required on-board estimates.

This plan was then refined over several versions, each of which was developed in a few days using this iterative process. The resource information and other details in the plan enabled mission designers to make informed decisions on how to shape the mission.

~~(a mission of 18 days, six days shorter than for MAMM) took over a year.~~

In planning the mission, ASPEN was used for several preliminary analyses and what-if scenarios. Specifically,

1. Determine the resource requirements for purposes of costing the mission and negotiating spacecraft resource allocations with the CSA.~~What are the usage statistics for the spacecraft instruments (e.g., total SAR on-time, total number of data takes)?~~
2. How do different downlink scheduling policies impact the mission plan?~~How should the priorities be structured for the downlink stations (e.g., should downlinks go to the Alaskan SAR Facility before Gatineau SAR Station when there are downlink opportunities at both)?~~
3. What is the impact of not using certain ground stations?~~Should certain ground stations be used at all during the mission (e.g., can we drop Svalbard Station)?~~

The ASPEN system answered each of these questions.

Question (1) was addressed using a summary page containing usage statistics generated after each analysis run of ASPEN, which gave the vital statistics for on-board recorder usage, SAR on-time, and total downlink data time broken down by station. This ~~helped in~~was invaluable in negotiati~~o~~n~~s~~g on-board resource allocations~~ at very preliminary stages of the mission to show that the coverage was capable within the time constraints and spacecraft limitations, based on very preliminary swath sets and an ASPEN-generated downlink schedule~~. The detailed downlink schedules were used to determine ground station cost estimates, and schedule resources. The early and detailed availability of these schedules greatly simplified this process over AMM-1.

Question (2) was addressed by performing what-if simulations using the ASPEN system. Since downlink station priorities were one of the parameters of the downlink generation algorithm, the plan was expanded and downlink schedules generated using four different possible priority systems, based on the actual cost to downlink data to certain ground stations. ASPEN supplied the data to reach a decision on the priorities and significantly impact the mission negotiations during the early stages.

Question (3) was addressed using similar what-if scenarios, where ASPEN was restricted from creating downlinks to certain stations. This enabled a closer examination of the impact of removing a ground station on the other stations and on the science collection in general. Using this information, the mission eliminated an unnecessary ground station early on in the mission operations planning phase, and saved a significant amount of funding ~~which would have gone to setting up and maintaining communication with that station~~that would have been needed to support that station during operations.

During the mission

Data takes missed due to spacecraft or ground station anomalies during the first cycle can be rescheduled for later in the cycle. The automated planner was available during operations for identifying operations conflicts in manually generated replan schedules. The system took as input ~~In order to test the system for operations, several informal tests were conducted using the ASPEN system. Given a potential~~ the replanned schedule~~et of swaths and other required files (such as a downlink plan without certain station masks where there were station outages)~~, and provided a list of conflicts~~ASPEN could produce the list of conflicts in minutes~~ within minutes. This capability enabled the ~~is would quickly allow the~~ replanning team to quickly ~~see~~ identify and correct any constraint violations before ~~if their plans satisfied the mission and spacecraft constraints before~~ submitting it to the MMO for a ~~second~~ final (and more costly) check.

These capabilities were demonstrated during operations rehearsals, and were available for use in operations, but were not in fact needed during the mission. Few anomalies occurred ~~i~~In the ~~initial 24 days of the mission~~first cycle, and they impacted swaths that could be trivially and confidently rescheduled manually.~~, there were no anomalies, which caused data loss that could not be trivially replanned by the on-site replanning team using the SPA tool, so we did not get to use ASPEN during the actual mission.~~

~~6.~~5. FUTURE ~~PLANS~~ WORK~~AND CONCLUSIONS~~

Although replanning was trivial on MAMM, it was much more difficult on AMM-1. In general, the ability to quickly and automatically reschedule observations lost to anomalies would be beneficial to many missions.

Anomalies during AMM-1 operations resulted in the loss of several data acquisitions. AMM-1 suffered 10 satellite anomalies and lost a primary ground receiving station. When anomalies occur, the missed observations must be rescheduled to meet the science requirements. The operations re-planning staff had between 48 and 72 hours in which to find up to 30 replacement swaths and submit a new schedule. To manually turn around plans within these time constraints required a team of four people working from pre-generated contingency plan segments. The missed observations were placed into gaps in the original plan to minimize coverage holes. More extensive changes, such as altering the remaining (unexecuted) planned swaths were avoided to minimize the planning effort and the chance of introducing errors into the plan. Unfortunately, it was sometimes impossible to find a way to reschedule all the missed observations within that time frame using these manual procedures. These observations were simply dropped from the schedule. Automated replanning during operations would allow faster turn-around with fewer people, and enable more extensive changes to the schedule in order to maximize science return.

Automated replanning could be implemented as follows. The user provides a set of swath opportunities for each missed observation, based on possibilities generated using the swath analysis tools. ASPEN selects at most one opportunity for each observation such that the resulting plan satisfies the operational constraints and recovers as many missed observations as possible. In selecting these observations, it avoids introducing conflicts and/ resolves them by modifying the schedule. The ASPEN search facility is designed for solving these kinds of problems.~~Based on this success, ASPEN will be used in a number of upcoming missions for earth-looking radar satellites. For example, a study on the LightSAR mission at NASA enabled the development of a similar model for full SAR coverage of Greenland using the LightSAR satellite, in order to perform what-if scenarios with the design and orbit of the spacecraft in order to increase science return while keeping the costs down. Furthermore, we are working with the Alaskan SAR Facility to use ASPEN on a day-to-day basis at their facility with all of the satellites they are currently involved. This would involve reconfiguring the current ASPEN model so that it reflected the constraints of alternative satellites. One future goal may be to parameterize such a model so that the system could be used for planning the behavior of any satellite by changing a fixed number of constraints.~~

~~More abstractly, the problem of finding an optimal set of swaths to achieve complete coverage of an area, given a tightly constrained spacecraft is a difficult problem, which currently does not have a fast, provably optimal solution that can be used for automated planning. We are working on this and other related algorithms to incorporate in the planner as a special-purpose algorithm to increase the robustness and ease of planning and scheduling SAR missions.~~

6. CONCLUSIONS

Automated planning created a significant savings in developing mission plans, and optimized science return in a way that manual planning would take too long to perform. The planning systems also enabled rapid generation of "what-if" plans for feasibility studies, mission costing, and resource negotiations. These studies directly contributed to the quality and success of the mission, and the mission planners considered this capability an invaluable tool. Automated planning was overwhelmingly successful for MAMM, and we would expect similar successes for future missions that employ this technology.

7. ACKNOWLEDGEMENTS

~~7. ACKNOWLEDGEMENTS~~

This paper describes work performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract from the National Aeronautics and Space Administration. The authors also wish to acknowledge the support of the Canadian Space Agency in providing the RadarSAT operations constraints, and the support and sponsorship of Dr. Kim Partington during his tenure as manager of NASA's Polar Science Program.

REFERENCES

[1] Chien, S., Rabideau, G., Knight, R., Sherwood, R., Engelhardt, B., Nutz, D., Estlin, T., Smith, B., Fisher, F., Barrett, T., Stebbins, G., Tran, D. "ASPEN – Automating Space Mission Operations using Automated Planning and Scheduling," *SpaceOps 2000*, Toulouse, France, June2000.

[2] Fukunaga, A.; Rabideau, G.; Chien, S.; and Yan, D. "Toward an Application Framework for Automated Planning and Scheduling," in *Proceedings of the 1997 International Symposium on Artificial Intelligence, Robotics, and Automation for Space*, Tokyo, Japan (1997).

[3] Garey and Johnson. *Computers and Intractability: A guide to the theory of NP-completeness*. W.H. Freeman and co., New York (1979).

[4] Rabideau, Knight, Chien, Fukunaga, and Govindjee, "Iterative Repair Planning for Spacecraft Operations in the ASPEN System" *International Symposium on Artificial Intelligence, Robotics, and Automation for Space (ISAIRAS)*, Noordjwijk, The Netherlands, June 1999.

[5] Sherwood et al, "ASPEN User's Guide" Internal Technical Document D-15482, Jet Propulsion Laboratory, California Institute of Technology. April, 1999. Also available at http://aspen.jpl.nasa.gov

[6] Zweben, M.; Daun, B.; Davis, E.; and Deale M., "Scheduling and Rescheduling with Iterative Repair," in Intelligent Scheduling, Morgan Kaufman, San Francisco (1994).

[7] ~~WEB:~~ h~~H~~ttp://www.radarsolutions.dera.gov.uk/swath.htm

[8] ~~Web:~~ http://www.stk.com

BIOGRAPHIES

Ben Smith is a Senior M~~m~~ember of ~~—~~the Exploration Systems Autonomy Section ~~technical~~ ~~staff in the Artificial Intelligence Group~~ at JPL. He led the development of the MAMM Automated Planning System, and is currently pursuing research in planning and scheduling for space applications, and in verification of autonomous systems . ~~Ben's research interests are in Machine Learning, Autonomous Agents, and Planning and Scheduling.~~ His other re-search interests include intelligent agents and machine learning. ~~He is the JPL lead of the planning component of the Remote Agent Experiment, which will autonomously control the Deep Space One spacecraft for several days. He also supports other AI group projects in autonomy and planning. Dr. Smith~~He received his Ph.D. from the ~~Computer Science department at the~~ University of Southern California ~~Information Sciences Institute~~ in ~~December~~ 1995. ~~Dr. Smith received his M.S. in computer science from USC in 1990 and his B.S. from the University of the Pacific in 1988.~~

Barbara Engelhardt is a member of the technical staff in the Artificial Intelligence Group at JPL, and a member of the Replanning Team for MAMM during the mission. Barbara's research interests are in machine learning, specifically learning spacecraft behavior from data, and planning and scheduling. Barbara received her M.S. in Computer Science from Stanford University working on Bayesian Learning Techniques, and her B.S. from Stanford University, both in 1999.

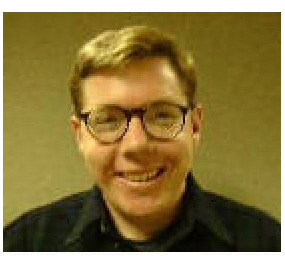

Darren Mutz is a member of the technical staff in the Artificial Intelligence Group at JPL. Darren received the Bachelor of Science degree in Computer Science at the University of California, Santa Barbara in 1997. His other current projects include Long Range Science Rover (LRSR), statistical hypothesis evaluation, and ASPEN.

John Crawford is the JPL Project Manager for the Modified Antarctic Mapping Mission and the Mission Architect for the Antarctic Mapping Mission in 1997. ~~He supervised the development of the specifications for the ASPEN model and gave critical user feedback on using the tool before and during the mission.~~ Furthermore, he has a point somewhere ~~somewhere~~ in Greenland named after him.

~~him.~~